Diagnosis of Complex Dynamical Systems

Many systems relevant to the Exploration Vision, for example planetary rovers, are unsuited for traditional approaches to diagnosis. Here we describe advanced techniques that extend diagnosis capabilities to these systems.

Background

Diagnosis is an important technology for many of the vehicles and systems relevant to the Exploration Vision. For autonomous systems such as planetary rovers, it provides state estimates that are used to decide appropriate courses of action. For future space and lunar habitats it can allow crew to spend far less time finding and repairing faults by identifying them automatically, resulting in much greater crew productivity than on, for example, the International Space Station. Diagnosis can also provide estimates of component performance or wear that can be used to reduce vehicle turnaround times by providing prognostic information about which subsystems to replace or repair, and which should continue to function normally.

Traditional approaches to diagnosis, such as Livingstone, build a discrete model of the system being diagnosed, and compare predictions from that model with discretized sensor inputs to determine when the behavior of the system does not match that of the model. These 'conflicts' are then used to search for other model states (usually faults) that do match system behavior correctly. For many systems, particularly digital systems, this approach is very effective, and Livingstone has been applied successfully in many domains, including as part of the Remote Agent Experiment that flew on the Deep Space 1 mission.

Unfortunately, the discrete approach of Livingstone is unsuitable for a large class of systems that cannot be easily discretized. An example of such a system is a planetary rover, in which the interaction between the rover and its environment makes discretization almost impossible. Another is life-support systems, or almost any system where quantitative information is important. We represent these systems quantitatively, typically with differential equation models, and have developed appropriate tools for tracking both the continuous state of these systems, and their discrete system modes (typically faults).

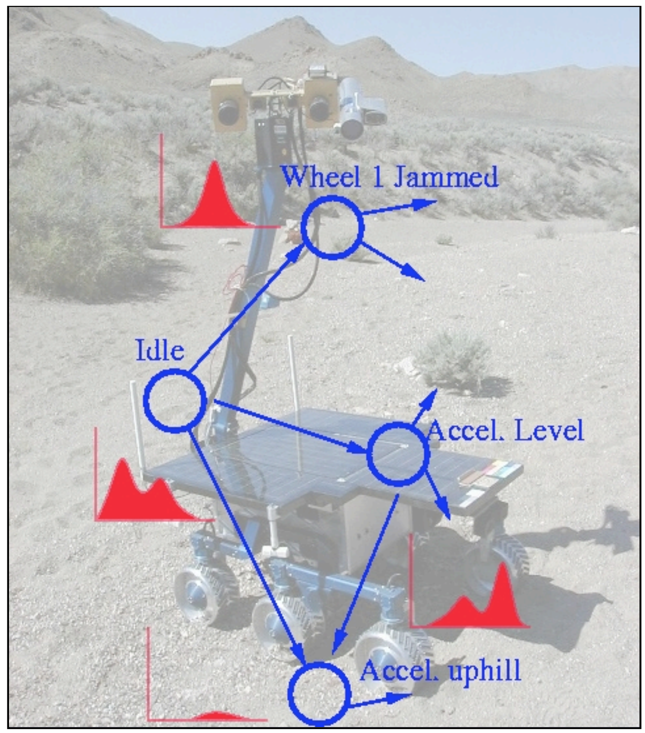

Research Overview

The diagnosis problem can essentially be characterized as follows: Given a model of the system to be diagnosed, and a stream of sensor information from the system, maintain an estimate of the state or states that the system could be in based on all the telemetry seen so far. The approach used in this technology is to construct a *hybrid* system model—one that represents the system as a set of discrete operating modes (corresponding to control modes and any faults that may occur), as well as a set of continuous system parameters. The continuous behavior of the system is described by differential equations that will typically be different for each discrete mode. Diagnosis with such a model consists of estimating both the discrete mode that the system is in and the values of the continuous parameters.

Diagnosis of Complex Dynamical Systems

The other important difference between this approach and traditional diagnosis technologies is the use of probabilistic diagnoses. Rather than maintaining a single candidate system diagnosis, or a small set of possible diagnoses, we provide a probability distribution over the possible states of the system. This is advantageous for a number of reasons: First, the probability distribution summarizes all the telemetry received by the diagnosis engine so far, making it easier to detect the onset of gradual faults, or faults that are initially hard to distinguish from nominal system behavior. Secondly, this probabilistic behavior can be of use in decision making by allowing the effects of planned future actions to be evaluated in states that have low probability but potentially catastrophic outcomes, rather than only in the most likely state. Finally, probabilistic information is of use for prognostics and maintenance, providing information about components that could have failed, or are degraded but not yet faulty.

The approach we take to producing probabilistic diagnoses of hybrid systems is a form of Bayesian belief updating—we maintain a probability distribution that represents our belief about the state of the system, given all the telemetry seen so far, and update that distribution whenever new telemetry arrives to reflect the new evidence. Because doing this exactly is intractable for the kinds of systems we are interested in, we approximate it using *particle filtering* algorithms.

Particle filtering is a Monte Carlo approach to state estimation that is extensively used in signal processing, and more recently in robot localization algorithms such as SLAM. It works by representing the probability distribution over the states of the system as a set of samples, each of which represents a possible state of the complete system. When new telemetry arrives, each sample randomly predicts a possible future state, and is then compared with the observation. Samples which are inconsistent with the observation are discarded, and are replaced at the next step by samples that do predict the observation well.

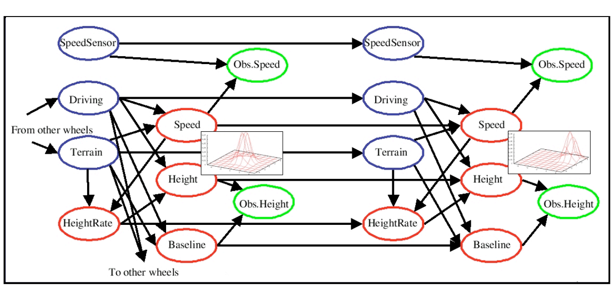

We have developed a number of innovative extensions to the basic particle filtering algorithm specifically to address the special characteristics of diagnosis problems, and to make the algorithms suitable for running on-board typical space platforms with their limited computational and memory resources. We have demonstrated the approach on the K-9 rover test bed and shown that we can accurately track the discrete state of the system in extremely noisy data, and can also reliably estimate continuous parameters of the system, even when those parameters are gradual system faults that cannot be directly observed.

Relevance to Exploration Systems

The Exploration Vision includes many complex vehicles and systems that can benefit from this technology. This includes habitats, which require diagnosis of their life-support and other systems, many of which require quantitative reasoning to diagnose correctly, and reusable launch vehicles which can greatly benefit from diagnosis techniques that can track gradual degradation of systems and components to not only detect faults when they occur, but also to predict future faults to make maintenance more efficient and reduce vehicle turnaround times.

H&RT Program Elements:

This research capability supports the following H&RT program /elements:

ASTP/Software, Intelligent Systems & Modeling

Points of Contact:

Richard Dearden
650-604-5616; dearden@email.arc.nasa.gov
http://ic.arc.nasa.gov/tech/group.php?gid=7&ta=2

Sriram Narasimhan
650-604-0832; sriram@email.arc.nasa.gov